

FINAL REPORT

UXO Discrimination Study Former Spencer Artillery Range

SERDP Project MR-1708

APRIL 2013

Larry Carip
Levi Kennedy
Todd Jobe
Signal Innovations Group, Inc.

This document has been cleared for public release



This report was prepared under contract to the Department of Defense Strategic Environmental Research and Development Program (SERDP). The publication of this report does not indicate endorsement by the Department of Defense, nor should the contents be construed as reflecting the official policy or position of the Department of Defense. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the Department of Defense.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 04-24-2013		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 06/2012-05/2013	
4. TITLE AND SUBTITLE 2012 UXO Discrimination Study Former Spencer Artillery Range, TN, Final Report				5a. CONTRACT NUMBER MR-1708	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Larry Carin (PI), Levi Kennedy, Todd Jobe				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Signal Innovations Group, Inc. 4721 Emperor Blvd., Suite 330 Durham, NC 27703				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The main technical objective of the Former Spencer Artillery Range demonstration is to validate and substantially automate the SIG learning process using next-generation electromagnetic induction (EMI) sensor data for discriminating targets-of-interest. This process includes three major components, feature extraction, site learning and excavation. At Camp Spencer SIG tested three methods for generating dig lists that vary by the amount of training data required. The SIG discrimination process has proven effective at providing efficient site discrimination with minimal training data. Former Spencer Artillery Range provides another example of this. The addition of adaptive modeling approaches that use an a priori metric of site complexity has increased effectiveness of this approach.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Levi Kennedy
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (include area code) 919-323-3456

Table of Contents

Table of Contents	i
List of Figures	ii
List of Tables	ii
List of Acronyms	iii
Keywords	iii
Abstract	1
Objective	3
Background	4
SIG Discrimination Process Overview	4
Materials and Methods.....	6
EMI Multi-Dipole Model and Feature Extraction	7
Feature Selection with BENet.....	8
Semi-Supervised Classification	9
Non-Myopic Active Learning.....	10
Results and Discussion	11
Aggressive Approach.....	11
NAEVA MetalMapper – Open and Dynamic Areas	11
URS MetalMapper – Open and Dynamic Areas.....	13
Intermediate Approach.....	14
NAEVA MetalMapper – Open and Dynamic Areas	14
URS MetalMapper – Open and Dynamic Areas.....	15
Conservative Approach.....	16
NAEVA MetalMapper – Open and Dynamic Areas	16
URS MetalMapper – Open and Dynamic Areas.....	18
UXO Type Assignment.....	18
Missed UXO	19
Comparison between aggressive, intermediate and conservative approaches.....	22
Evaluation of site complexity	24
Conclusions and Implications for Future Research/Implementation.....	25
Literature Cited	25

Appendices.....	27
Appendix A: Points of Contact	27

List of Figures

Figure 1. Flow diagram of the SIG discrimination process.	6
Figure 2: Example plot of the polarizability of a UXO as a function of time (right). The left image is the UXO with the polarizability axes shown. Because of the distinctive shape of UXO, at least of two of the polarizability axes will be coincident; in this case M2 (red) and M3 (blue) are coincident.	7
Figure 3.. Receiver operating characteristic (ROC) curves for UXO classifier at SLO site with features selection using the BENet algorithm (red line) and without feature selection (blue line). The number of false alarms is lower for the classifier where feature selection was used.	8
Figure 4: A comparison between supervised and semi-supervised classifiers for a 2. Labeled data from both classes (red and green circles) are shown, along with unlabeled data (black dots). The supervised classifier is trained on only the labeled data and the decision boundary is shown (dotted line). The semi-supervised classifier is trained on both the labeled and unlabeled data and the decision boundary (solid line) makes the two classes linearly separable.	9
Figure 5. ROC curves for the Aggressive approach. NAEVA and URS sensors (top,bottom), Open and Dynamic areas (left, right).....	13
Figure 6. ROC curves for the Intermediate approach, NAEVA and URS sensors(top, bottom), Open and Dynamic areas (left, right).....	15
Figure 7. ROC curves for the Conservative approach, NAEVA and URS sensors(top, bottom), Open and Dynamic areas (left, right).....	17
Figure 8. UXO type cross-tabulation for the Open site.	19
Figure 9. UXO type cross-tabulation for the Dynamic site.	19
Figure 10. Polarizabilities for missed UXO in the NAEVA dig lists. Responses for NAEVA (left) and URS (right) are both shown.	20
Figure 11. Polarizabilities for missed UXO in the URS dig lists. Responses for NAEVA (left) and URS (right) are both shown.	22

List of Tables

Table 1. Program Office Performance Objectives for Discrimination Analysis	3
Table 2. Summary of the performance for the Aggressive (Aggr.), Intermediate (Inter.), and Conservative (Cons.) approaches. Values are counts of flagged anomalies. For training data, missed seeds are shown in parentheses since they were required to be part of the training data but were not part of the training in the classification models.	23
Table 3. Complexity metric for the current site and three former sites. Larger values suggest more intrinsic complexity and a more conservative modeling approach.	24

List of Acronyms

BENet	Bayesian Elastic Net
EMI	Electromagnetic Induction
ESTCP	Environmental Security Technology Certification Program
ISO	Industry Standard Object
MAP	Maximum <i>a posteriori</i>
MTL	Multi-task learning
NAEVA	NAEVA MetalMapper Data (Open and Dynamic areas)
nFA	Number of false alarms
NMAL	Non-Myopic Active Learning
PNBC	Parameterized Neighborhood-Based Classifier
QC	Quality Control
ROC	Receiver Operating Characteristic
Rx	Receiver
RVM	Relevance Vector Machine
SIG	Signal Innovations Group, Inc.
SLO	San Luis Obispo Demonstration Site
Spencer	Former Spencer Artillery Range
SVD	Singular Value Decomposition
TEMTADS	Time-domain Electromagnetic Multi-sensor Tower Array Detection System
Tx	Transmitter
URS	URS MetalMapper Data (Open and Dynamic Areas)
UXO	Unexploded Ordnance

Keywords

UXO, Bayesian classifier, active learning

Abstract

Objective: The main technical objective of the Former Spencer Artillery Range demonstration is to validate and substantially automate the SIG learning process using next-generation electromagnetic induction (EMI) sensor data for discriminating targets-of-interest. This process includes three major components: feature extraction, site learning and excavation. The end result of the discrimination is a list of classifications for all anomalies at a site. Discrimination performance is increased by maximizing the number of unexploded ordnance (UXO) correctly identified, maximizing the number of non-UXO anomalies correctly identified, specifying an appropriate dig no-dig boundary, and minimizing the number of anomalies that cannot be analyzed. Discriminations performed at Camp Spencer explored an adaptive approach to site learning that varies the discrimination approach and amount of training data depending on the expected difficulty of discriminating the site. SIG also proposes a metric for determination of site complexity to determine the correct approach to use based on the complexity of the site. More complicated sites require more aggressive discrimination methods and more training data to minimize the potential for missed UXO.

Technical Approach: SIG performed discriminations on datasets from two companies across two sites: Open and Dynamic. Both companies, NAEVA and URS, used the MetalMapper sensor. For each company and site a set of three discrimination methods. These discrimination methods were designed to accommodate sites that had differing levels of intrinsic complexity and available information about the types of UXO present. These methods, labeled aggressive, intermediate and conservative, varied the amount of training data requested from 0 to around 50 labels. Classification for the conservative method utilized a generative Bayesian classifier while the intermediate and aggressive methods used a semi-supervised, parametric Bayesian classifier. To compliment these discrimination methods a complexity metric was developed that allows for *a priori* selection among the three modeling approaches for a new site.

Results: Summarized by modeling approach.

Aggressive: The aggressive approach acquired UXO with the fewest clutter dug, but tended to miss UXO types that had no examples in the test pit data. No training data were acquired, and only test pit data were used for generative model. None of the 23 UXO were missed in the Dynamic area with approximately 60 digs for the NAEVA and URS datasets. Three of the 86 UXO were missed in the Open area with the NAEVA data and two UXO were missed with URS data having 213 and 300 digs, respectively. All the missed UXO were due to either selecting a one-anomaly inversion model where a two-anomaly inversion was appropriate or was of a UXO type not present in the test pit data.

Intermediate: This approach required the most digs to capture UXO, primarily because the stop-dig threshold had to be altered significantly in order to capture missed QC seeds which were fit to one-anomaly models, but were actually two-anomaly digs. 13 training labels and 3 training labels were acquired for the NAEVA and URS datasets, respectively. No UXO were missed in

the Dynamic area with approximately 80 digs for the NAEVA and URS datasets. One UXO was missed in the Open area for both the NAEVA and URS datasets out of 387 and 363 digs.

Conservative: This approach was the only one to capture all UXO in both the Open and Dynamic areas. The NAEVA dataset and the URS dataset relied upon 42 and 47 actively learned training labels. The NAEVA dataset missed 12 QC seeds, while the URS dataset missed only 2. The final NAEVA classification missed no UXO with 369 digs, while the URS classification missed one UXO with 291 digs.

Complexity analysis of former sites Pole Mountain, Camp Spencer and Camp Beale along with the current site suggest that Camp Spencer would be best discriminated using the intermediate or conservative approaches. Results bear this out in that fewer UXO were missed with the intermediate and conservative approaches, than the aggressive approach.

Benefits:

The SIG discrimination process has proven effective at providing efficient site discrimination with minimal training data. Former Spencer Artillery Range provides another example of this. The addition of adaptive modeling approaches that use an *a priori* metric of site complexity has increased effectiveness of this approach.

Objective

The main technical objective of the SIG Former Spencer Artillery Range (Spencer) UXO discrimination demonstration is to validate and substantially automate the SIG learning process using next-generation electromagnetic induction (EMI) sensor data for discriminating targets-of-interest. All elements of human interpretation and intuition are being incrementally constrained or removed from the process, resulting in an automated process, where all algorithm parameters and thresholds will either be determined by specified site parameters (i.e., expected or inferred munitions types) or by data-driven inferences (i.e., cross-validated operating threshold). SIG applied and matured each of the three key process phases that constitutes the SIG statistical learning approach to UXO discrimination. The three phases of discrimination process include: Phase I - feature extraction, Phase II – site learning, and Phase III – excavation. Each of the phases is described in detail below. Validation of discrimination process entails meeting all of the discrimination performance objectives defined by the program office for each of the sites considered (Table 1). In particular Phase II has been examined in-depth at Camp Spencer. A

Table 1. Program Office Performance Objectives for Discrimination Analysis

Performance Objective	Metric	Data Required	Success Criteria
Analysis and Classification Objectives			
Maximize correct of classification of targets of interest	Number of targets-of-interest retained.	<ul style="list-style-type: none"> • Prioritized anomaly lists • Scoring reports from the Institute for Defense Analyses (IDA) 	Approach correctly classifies all targets-of-interest
Maximize correct classification of non-target of interest (UXO)	Number of false alarms eliminated.	<ul style="list-style-type: none"> • Prioritized anomaly lists • Scoring reports from IDA 	Reduction of false alarms by > 65% while retaining all targets of interest
Specification of no-dig threshold	Probability of correct classification and number of false alarms at demonstrator operating point.	<ul style="list-style-type: none"> • Demonstrator - specified threshold • Scoring reports from IDA 	Threshold specified by the demonstrator to achieve criteria above
Minimize number of anomalies that cannot be analyzed	Number of anomalies that must be classified as “Unable to Analyze.”	<ul style="list-style-type: none"> • Demonstrator target parameters 	Reliable target parameters can be estimated for > 98% of anomalies on each sensor’s detection list.

site adaptive discrimination approach was developed and tested. This approach uses one of three models labeled: aggressive, intermediate and conservative. The choice of model depends on an *a priori* measure of site complexity. Each of the three modeling approaches was applied to Camp Spencer and then the results were interpreted in light of the calculated complexity of the site.

Background

Using next-generation cued sensors, discrimination performance on real sites has shown the feasibility of advanced statistical analyses for distinguishing unexploded ordnance (UXO) from clutter. Signal Innovations Group, Inc. (SIG) has demonstrated the effectiveness of site-specific statistical learning for smartly selecting labeled training data to maximize target discrimination. This technology has been developed and validated under previous SERDP efforts by SIG and Duke University and was ready for application to ESTCP site demonstration at Former Spencer Artillery Range.

Many current analysis approaches rely on expert scientists to make educated decisions at multiple points in the discrimination analysis process. This situation is not scalable, transferable, or cost effective. The SIG approach standardizes the options and creates a documented process flow that can be explicitly followed. Key statistical technologies that were validated and automated during this effort include: physics-based target/sensor models, subspace denoising, automated and efficient feature extraction, data selection for classifier training. These techniques represent the state of the art in digital geophysics.

SIG Discrimination Process Overview

The SIG discrimination process laid out in [5] can be summarized in the following ‘recipe’ (Figure 5):

- **Data Conditioning** - First, raw, unlabeled anomaly data are received. Then, quality control checks are performed. These include ensuring that the background subtractions are complete and determining if the raw sensor inputs need any scaling in order to be appropriate for the feature extraction software.
- **Subspace Denoising** - The anomaly data is denoised to ensure robust performance for discriminating late time-gate features.
- **Feature Extraction** - A robust multi-anomaly dipole model is fitted to the data. The polarizability parameters from this fitting become the set from which features are drawn for classifier training. In addition to the time-domain polarizabilities, a set of 9 ‘rate’ features were calculated. These features were the calculated by fitting the time-domain polarizabilities of each axis to an exponential-decay model:

$$p_i = r_{1i} + r_{2i} e^{\frac{-t}{r_{3i}}}$$

where $i \in \{x, y, z\}$ is the current axis, p is the polarizability, t is time and $\{r_1, r_2, r_3\}$ are the fitted rate parameters. Though r_{1i} is unphysical, it is useful for adjusting for noise at late time gates and where odd responses would make the optimization difficult. The optimized values of the rate parameters were found using non-linear least squares.

- **Basis Selection** - A few of the many possible features are selected based on their physical interpretation as they relate to the anomaly, and, using these features, the most informative set of anomalies are selected via an information metric to begin classifier training.
- **Feature Set Augmentation** - The feature set is then augmented by adding early, mid and late time polarizabilities values.
- **Automated Feature Selection** - For the now larger feature set, the most relevant set of features is selected using BENet.
- **Semi-supervised PNBC Training (STL or MTL)** - When the PNBC is trained only using data from the current site of interest, it is called Single Task Learning (STL). When the PNBC is trained for multiple sites simultaneously it is called multi-task learning (MTL).
- **Non-myopic Active Learning** - Based on the estimates made with the PNBC classifier, a new set of anomalies will be selected for labeling using NMAL. The goal at this step is to maximize the information gain from new labels requested from the set of unlabeled anomalies. The process is repeated as the PNBC classifier adequately learns data manifold. The stopping criteria for the learning process is apparent when the remaining unlabeled data points have approximately equal information for improving the classifier. At which point, labeling any one-anomaly is no better than any other.
- **Excavation Adapted Threshold Selection** - At this point, the highest probability UXO are selected for excavation and labels. The classifier continues to be retrained when new labels are revealed. This process continues until the highest probability UXO items excavated are all found to be clutter at which point digging stops.

The process outlined above falls into 3 broad phases: Feature Extraction, Site Learning, and Excavation. Details on each phase are given in the next subsections. The SIG Discrimination process is relatively linear save for two feedback steps. The first feedback is in training the semi-supervised classifier, where additional anomaly labels are requested until the classifier reaches sufficient stability. The second feedback is during the excavation of anomalies, where the classifier is retrained with additional labeled anomalies until either the UXO/clutter predictions become highly separable or until high probability anomalies are substantially revealed to be clutter upon excavation.

For the Camp Spencer discrimination the second feedback was eliminated in the current analysis to focus on the first feedback. Thus differences in performance among the different modeling approaches (aggressive, intermediate, conservative) relied solely on the differences in training data acquired during the site learning phase.

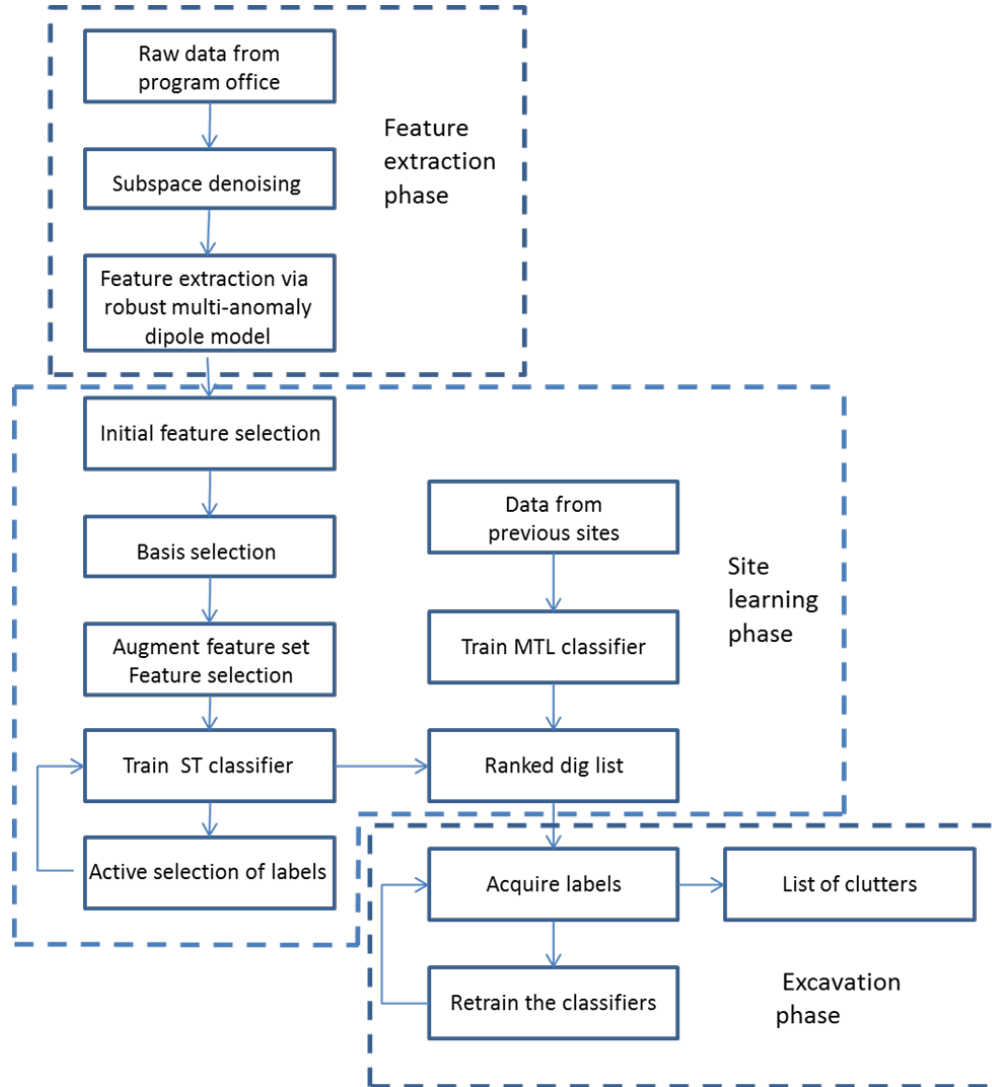


Figure 1. Flow diagram of the SIG discrimination process.

Materials and Methods

SIG applied the discrimination process at Camp Spencer for the MetalMapper sensor data collected by NAEVA and URS in the ‘Open’ and ‘Dynamic’ sites. For the purposes of training, the ‘Open’ and ‘Dynamic’ sites were treated jointly. The process for each dataset involved the following key technologies, including: parametric target/sensor modeling, robust feature extraction and selection, semi-supervised classifier training using active selection of labeled data and multi-task learning incorporating past demonstration site data. To assess the impact of various intensities of training, discrimination models were created that relied upon no training data, fewer than 20 training data points, and more than 20 training data points. These were labeled the aggressive, intermediate, and conservative approaches, respectively. Each of the aforementioned technologies is described briefly in the following subsections.

EMI Multi-Dipole Model and Feature Extraction

SIG extracted features by fitting raw sensor data to a physics-based parametric model [1], [2] that was developed under SERDP support and has been successfully demonstrated and validated in Camp Sibert, Camp San Luis Obispo (SLO), and Camp Butner analyses. It has been shown [3], [4] that the induction response of simple targets can be efficiently represented in terms of one or more time/frequency-dependent magnetic dipoles. In particular, the magnetic dipole moment m of a target is represented as $m = M \cdot H^{inc}$, where H^{inc} represents the incident (excitation) magnetic field and M is a tensor that relates H^{inc} to m . Using reciprocity in wave propagation, the total magnetic field observed at the receiver coil, H^{rec} , can be represented as

$$H^{rec}(\omega) \propto r_{st} \cdot U^t \cdot M \cdot U \cdot r_{st}$$

where r_{st} is a unit vector directed from the source to the target and the 3×3 unitary matrix U (that contains information about the target orientation) rotates the fields from the coordinate system of the sensor to the coordinate system of the target. To simplify the above expression, an

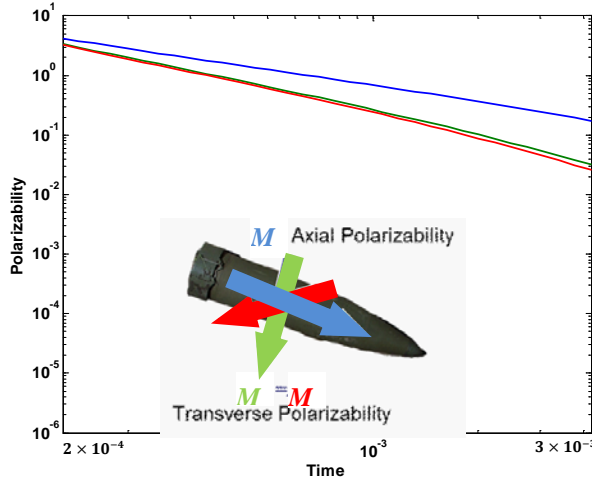


Figure 2: Example plot of the polarizability of a UXO as a function of time in seconds. Both axes are log scaled. The inset image is the UXO with the polarizability axes shown. Because of the distinctive shape of UXO, at least of two of the polarizability axes will be coincident; in this case M2 (green) and M3 (red) are coincident.

assumption is made that the source responsible for the incident field H^{inc} can be characterized by a dipole. This assumption is valid if the dimension of the transmitter coil is much smaller than its distance from any buried anomaly. While this may be appropriate for many sensors, it is not appropriate for the multi-coil multi-axis sensors (e.g. MetalMapper). SIG has developed a model to synthesize H^{inc} for these systems; wherein the physical sizes/shapes of the transmitter and receiver coils are accounted for explicitly (this is done with a rigorous Biot-Savart analysis). This model has already been validated on MetalMapper data (under SERDP contract MR-1708).

Based on the generalized forward model described above, the three principal polarization terms in the magnetic polarization dyadic, the orientation angles in the rotation matrix, as well as the position of the object can be extracted directly from the field data by using a nonlinear least-square solver (Figure 1). It is well known that the trigonometric functions in the rotation matrix are nonlinear multi-valued functions that result in many local solutions. However, the tensor is a symmetric matrix, having only six independent elements (parameters), and the general magnetic polarization dyadic is a linear function of the measured field data. This fact is exploited by extracting the six parameters directly from the measured data (using least-square inversion), rather than extracting the three principal polarization components and three rotation angles directly. The number of non-linear parameters in the model is greatly reduced in this implementation, and the problem of local solutions is significantly relieved, resulting in reliable

convergence of the feature inversion. This procedure is also capable of performing simultaneous feature inversion for multiple co-located (or nearby) anomalies, which was validated during the SIG SLO demonstration.

Feature Selection with BENet

Adaptive learning of a classifier *in situ* benefits from refining the appropriate set of extracted features for the targets under test. This occurs because of the ‘curse of dimensionality’ where the number of data points required to cover the breadth of a features space grows exponentially with the number of features considered. If the amount of training data does not sufficiently sample the feature space, then the learned classifier will lack statistical support and class estimate uncertainty is large. At SLO in particular, feature selection played a key role in classifier performance (Figure 3). Bayesian classification models perform feature selection by




Figure 3. Receiver operating characteristic (ROC) curves for UXO classifier at SLO site with features selection using the BENet algorithm (red line) and without feature selection (blue line). The number of false alarms (nFA) is lower for the classifier where feature selection was used.

placing a sparseness prior on the inferred feature weights. The Bayesian elastic net (BENet) regression model used for feature selection employs a sparseness prior equivalent to a convex combination of L1-norm and L2-norm penalties in a least squares optimization formulation [5], [2]. The sparseness prior of the BENet model jointly infers the essential subset of relevant features, including correlated features, for a given classification task. Rather than encouraging the selection of a single feature in a set of correlated important features (like similar approaches such as RVM), the BENet model encourages the selection of all correlated important features. By performing sparse and grouped feature selection, the BENet algorithm provides a more robust approach to

feature adaptability and the interpretation of important features, ultimately requiring fewer training data samples to achieve robust statistical support.

Site-Adaptive Model Complexity

A key result from previous discriminations at Pole Mountain, Camp Beale, Camp Butner, and Camp San Louis Obispo, is that sites vary in the quality of features extracted from anomalies. This may be due to differences among sensor technologies, differences in soil types, or differences in data collection methods. As a result, it is ‘easier’ to discriminate UXO from clutter at some sites than others. At easy sites simple models, such as a library matching approach, are effective and dig all UXO with very few clutter. More difficult sites may require a more complex model to account for increased uncertainty in the feature responses. The SIG approach to discrimination seeks to adaptively vary model complexity based on intrinsic properties of the unlabeled site data. This complexity can be measured *a priori* and the appropriate modeling approach selected before training begins.

Complexity of a site is measured using an information theoretic metric applied to the distribution of unlabeled polarizability features. Complexity of a data generating process, $C(P)$, is dependent on two factors: the information stored in the system, $H(P)$, and the disequilibrium of the system, $D(P)$ (i.e. the distance of the underlying process generating the observed data and a uniform distribution) [6] so that,

$$C(P) = H(P) \times D(P)$$

The system information is measured as the ratio of differential entropy, I , of the data to the entropy of a similarly sized uniform distribution, I_{max} : $H = I/I_{max}$. Disequilibrium is measured as the Kullback-Leibler divergence between the observed distribution and a uniform distribution with a similar range of features: $D = (I_{max} - I)/I_{max}$. Each feature in the observed data is transformed to a standard Gaussian distribution so that the differential entropy is defined for the observed data as:

$$I = \frac{1}{2} \ln((2\pi e^k) |\Sigma|)$$

where Σ is the sample covariance matrix with k features. The resulting complexity is distributed (0, 1) with high values representing high complexity, and low values representing low complexity.

When C is high for a given site, the expectation is that the site would be difficult to discriminate. A classification approach requiring more training data would be used (i.e. the ‘complex’ approach defined here). When C is low for a given site, the site should be relatively easy to discriminate and the ‘aggressive’ approach would be used. These relations describe the relative




Figure 4: A comparison between supervised and semi-supervised classifiers for a two-feature dataset. Labeled data from both classes (red and green circles) are shown, along with unlabeled data (black dots). The supervised classifier is trained on only the labeled data and the decision boundary is shown (dotted line). The semi-supervised classifier is trained on both the labeled and unlabeled data and the decision boundary (solid line) makes the two classes linearly separable.

values of C for a given modeling approach. The best modeling approach for a specific value of C must be estimated from a set of sites where all three classification approaches (‘complex’, ‘intermediate’, or ‘aggressive’) are applied. To facilitate this all three modeling approaches have been applied for the Camp Spencer datasets.

Semi-Supervised Classification

Semi-supervised learning is applicable to any sensing problem for which all of the labeled and unlabeled data are available at the same time, and therefore, particularly for the current demonstration study. Semi-supervised classification was employed in the intermediate and complex modeling approaches. In most practical applications (including the recent demonstration at Camp

Butner), semi-supervised learning has been found to yield superior performance relative to widely applied supervised algorithms. Figure 4 depicts the advantage of a semi-supervised approach to classification over its supervised counterpart. A classifier trained purely on labeled data (depicted as red and green circles) is shown as a purple dashed line and generates classification errors. In contrast, a semi-supervised classifier trained on both labeled and unlabeled data will generate perfect classification (depicted by the blue line). Note that the context provided by the unlabeled data was crucial in improving the classification performance in this case, since the labeled data were not representative of the two class distributions. As the number of training samples increases, the supervised classifier should approximate the semi-supervised classifier. Semi-supervised formulation treats the dataset (labeled and unlabeled) as a set of connected nodes, where the affinity w_{ij} between any two feature vectors (nodes) \mathbf{f}_i and \mathbf{f}_j is defined in terms of a radial basis function [7]. Based on the above formulation, one can design a Markov transition matrix $\mathbf{A} = [a_{ij}]_{N \times N}$ that represents the probability of transitioning from node \mathbf{f}_i to \mathbf{f}_j . Assuming $\mathcal{L} \subseteq \{1, 2, \dots, N_L\}$ represents the set of labeled data indices, the likelihood functional can be written as

$$(\{y_i, i \in \mathcal{L}\} | \mathcal{N}(\mathbf{f}_i), \boldsymbol{\theta}) = \prod_{i \in \mathcal{L}} p(y_i | \mathcal{N}(\mathbf{f}_i), \boldsymbol{\theta}) = \prod_{i \in \mathcal{L}} \sum_{j=1}^{N_i} a_{ij} p(y_i | \mathbf{f}_j, \boldsymbol{\theta})$$

where $\mathcal{N}(\mathbf{f})$ defines the neighborhood of \mathbf{f} . Estimation of classifier parameters $\boldsymbol{\theta}$ can be achieved by maximizing the log-likelihood via an Expectation-Maximization algorithm [8]. To enforce sparseness of $\boldsymbol{\theta}$ (enforcing most of the components of the parameter vector $\boldsymbol{\theta}$ to be zero), one may impose a zero-mean Gaussian prior on $\boldsymbol{\theta}$. A zero-mean Gaussian prior with appropriate variance can strongly bias the algorithm in choosing parameter weights that are most likely very small (close to zero). This technique is known as a parameterized neighborhood-based classifier (PNBC).

Non-Myopic Active Learning

Given that available training data labels at the beginning of a demonstration are not available and that excavations must be performed to reveal training data labels, one may ask in which order anomalies should be excavated to maximally improve the performance of the classifier algorithm. One useful criterion is to use the confidence on the estimated identity of the anomalies that are yet to be excavated. Specifically, one may ask which unlabeled anomaly label would be most informative to improve classifier performance if the associated label could be made available. It has been shown [9] that this question can be answered in a quantitative information-theoretic manner.

For active label selection, posterior distribution of the classifier is approximated as a Gaussian distribution centered on the maximum *a posteriori* (MAP) estimate. The uncertainty of the classifier is quantified in terms of the posterior precision matrix. The objective of AL is to choose a feature vector for labeling that maximizes the mutual information (I) between the classifier $\boldsymbol{\theta}$ and the new data point to be labeled. The mutual information can be quantified as the expected decrease of the entropy of $\boldsymbol{\theta}$ after new sample \mathbf{f}_{i^*} and its label y_{i^*} are observed.

$$I = \frac{1}{2} \log \frac{|H'|}{|H|} = \frac{1}{2} \log \{1 + p(y_{i*} | \mathbf{f}_{i*}, \boldsymbol{\theta}) \times [1 - p(y_{i*} | \mathbf{f}_{i*}, \boldsymbol{\theta})] \mathbf{f}_{i*}^T \mathbf{H}^{-1} \mathbf{f}_{i*}\}$$

It is important to note that the mutual information I is large when $p(y_{i*} | \mathbf{x}_{i*}, \boldsymbol{\theta}) \approx 0.5$. Hence, the AL prefers to acquire labels on those unlabeled samples for which the current classifier is most confused or uncertain. In this fashion the classifier learns quickly by not excavating anomalies that reveal redundant information. The process continues as new labels are revealed until the expected information gain for the remaining anomalies is approximately uniformly low. At that point the classifier is adequately trained and target inference on the remaining unlabeled anomalies can be reliably performed. By invoking the principle of submodularity in the algorithm optimization, the approach has been adapted to allow for the selection of multiple simultaneous labels at one time, making the technique operationally practical.

Results and Discussion

A main objective of the Camp Spencer discrimination was to differentiate between SIG's aggressive, intermediate and conservative approaches. These approaches differ in the amount of training data they required. Some of the initial dig lists based on these approaches missed quality control (QC) seeds, which were subsequently included in the dig lists as training data. To keep the amount of training data consistent between the different modeling approaches, however, the missed QC were not included in any of SIG's predictive models. Instead, the stop-dig thresholds were adjusted to capture the QC seeds.

The final dig list for every modeling approach and every sensor captured all UXO in the Dynamic area. Some approach/sensor combinations missed UXO in the Open area. The specific results for each modeling approach are presented below.

Aggressive Approach

In the aggressive approach, anomalies were classified without any training data save for the test pit data. A model was built for each UXO type in the test pit data for Camp Spencer by sensor. So, the NAEVA models were built on the NAEVA test pit data and the URS models were built on the URS test pit data. The features for these models were selected based on sparse Bayesian classifiers for previous sites. Given these features a generative model was created for each UXO type. This generative model provided probabilities of being a particular UXO type for each anomaly in the dataset. A stop-dig threshold was based on the minimum resubstituted probability of the test pit data and the mode of the predicted probability of the unlabeled anomalies.

NAEVA MetalMapper – Open and Dynamic Areas

The test pit data for the NAEVA dataset included UXO Types 105mm, 75mm, stokes mortar, medium Industry Standard Object (ISO), 37mm and small ISO. The small ISO in the test pit was apparently schedule 40 pipe with a diameter of 31mm. These were different than the small ISO labeled in the test data which were schedule 80 pipe with a diameter of 38mm. The features selected for each UXO type model were distinct and based on ground-truth from previous sites.

Given the set of appropriate features for a UXO type, a generative model was created based on the test pit responses only.

The initial dig list contained 213 digs for the Open site and 44 digs for the Dynamic site. These initial dig lists missed two Q.C. seeds (SR-1676, SR-1729) which were a small ISO (schedule 80 pipe) and a 37mm projectile. The small ISO was missed because there were no ISOs of that type in the test pit data. The aggressive approach is dependent upon representatives of the UXO types. The model based on the medium ISO test pit data, however, did allow for some match the schedule 80 pipe. So, there were fewer QC small ISOs missed in the NAEVA data than the URS data which did not have medium ISO test pit data. The 37mm projectile was missed because it should have been classified as a two-anomaly model, though the one-anomaly inversion model fit error was less than the cutoff SIG has used in the past to delineate one and two-anomaly models (0.05).

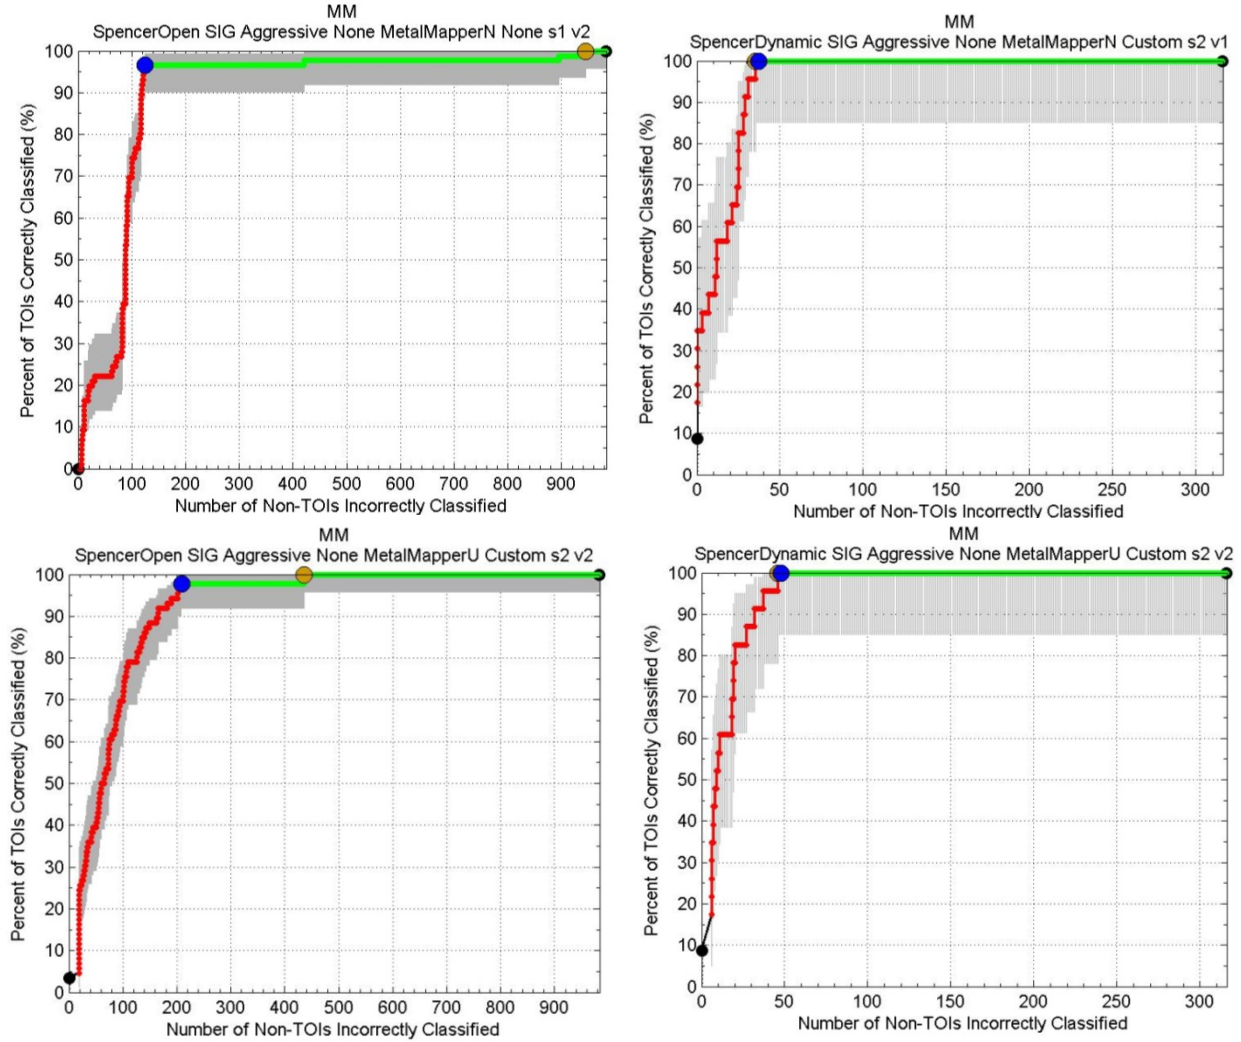


Figure 5. ROC curves for the Aggressive approach. NAEVA and URS sensors (top,bottom), Open and Dynamic areas (left, right)

Having adjusted the stop-dig threshold to accommodate the missed QC seeds, the final dig list for the Dynamic area captured all UXOs with 59 digs and the final dig list for the Open area remained 213 digs but missed 3 UXO (Figure 5, SR-181, SR-383, SR-710). Further details on the missed UXO are presented the ‘Missed UXO’ subsection.

URS MetalMapper – Open and Dynamic Areas

The test pit data for the URS dataset included UXO Types: 75mm, 37mm and small ISO. The small ISO in the test pit was apparently schedule 40 pipe with a diameter of 31mm. As with the NAEVA test pit data, these small ISO were different that the small ISO labeled in the field which were schedule 80 pipe with a diameter of 38mm. Building of the generative model followed similarly to the aggressive approach for the NAEVA dataset.

The initial dig list contained 220 digs for the Open site and 55 digs for the Dynamic site. These initial dig lists missed five Q.C. seeds (SR-199, SR-837, SR-873, SR-1502, SR-1676). All of the

missed seeds were attributable to either multi-anomaly responses that had good one-anomaly fits in SIGs inversion model, or were small ISOs that were not represented in the test pit data.

Having adjusted the stop-dig threshold to accommodate the missed QC seeds, the final dig list for the Dynamic area captured all UXOs with 71 digs and the final dig list for the Open area increased to 300 digs missing two UXO (Figure 5, SR-194, SR-633). Both of the missed anomalies were of the small ISO type not represented in the test pit data.

Intermediate Approach

The intermediate approach began from the same starting point as the aggressive one. That is, a set of generative models for each UXO type in the test pit data. But, the intermediate approach collected some training data based on these models. The active learning technique for acquiring training data was based on the minimum resubstituted probability of the test pit data. In short, a single training point was selected for each UXO type that was near this minimum resubstituted probability. If the revealed training label was UXO, this anomaly became the new minimum threshold and an additional training label was selected. If the revealed training label was clutter, digging stopped for that UXO type. Once a single clutter had been dug for all UXO types, training was completed. Discrimination was then performed as in the conservative approach using the PNBC.

NAEVA MetalMapper – Open and Dynamic Areas

The initial dig list contained 251 digs for the Open site and 82 digs for the Dynamic site. Of those initial digs, 13 were training. These initial dig lists missed nine Q.C. seeds (SR-572, SR-648, SR-790, SR-837, SR-873, SR-978, SR-1502, SR-1609, SR-1729). Half of the missed seeds were two-anomaly responses whose one-anomaly response had a low fit error so that the two-anomaly response was not selected. Upon learning that the missed QC seeds were in fact, two-anomaly, there were two options for how to proceed. The first option was to use the two anomaly model for the missed QC seeds and risk missing other two-anomaly UXO in the unlabeled data whose one-anomaly fits were also good. The second option was to try to capture any remaining two-anomaly UXO with good one-anomaly fits by training on the missed QC using their original one anomaly model. The latter option was chosen as it represented the more conservative of the two approaches and was consistent with the SIG model selection heuristic which favors selecting the one-anomaly model if its fit was good. The other missed QC seeds were near a clutter that was dug during training in feature space. Since there were so few training point, and a discriminative classifier was being used, these clutter had undue weight placed on them. The solution for such cases was to increase the variance of the non-linear kernels used for classification so that the bulk of the test pit data and training UXO obtained more influence than the single clutter dug in that area of feature space.

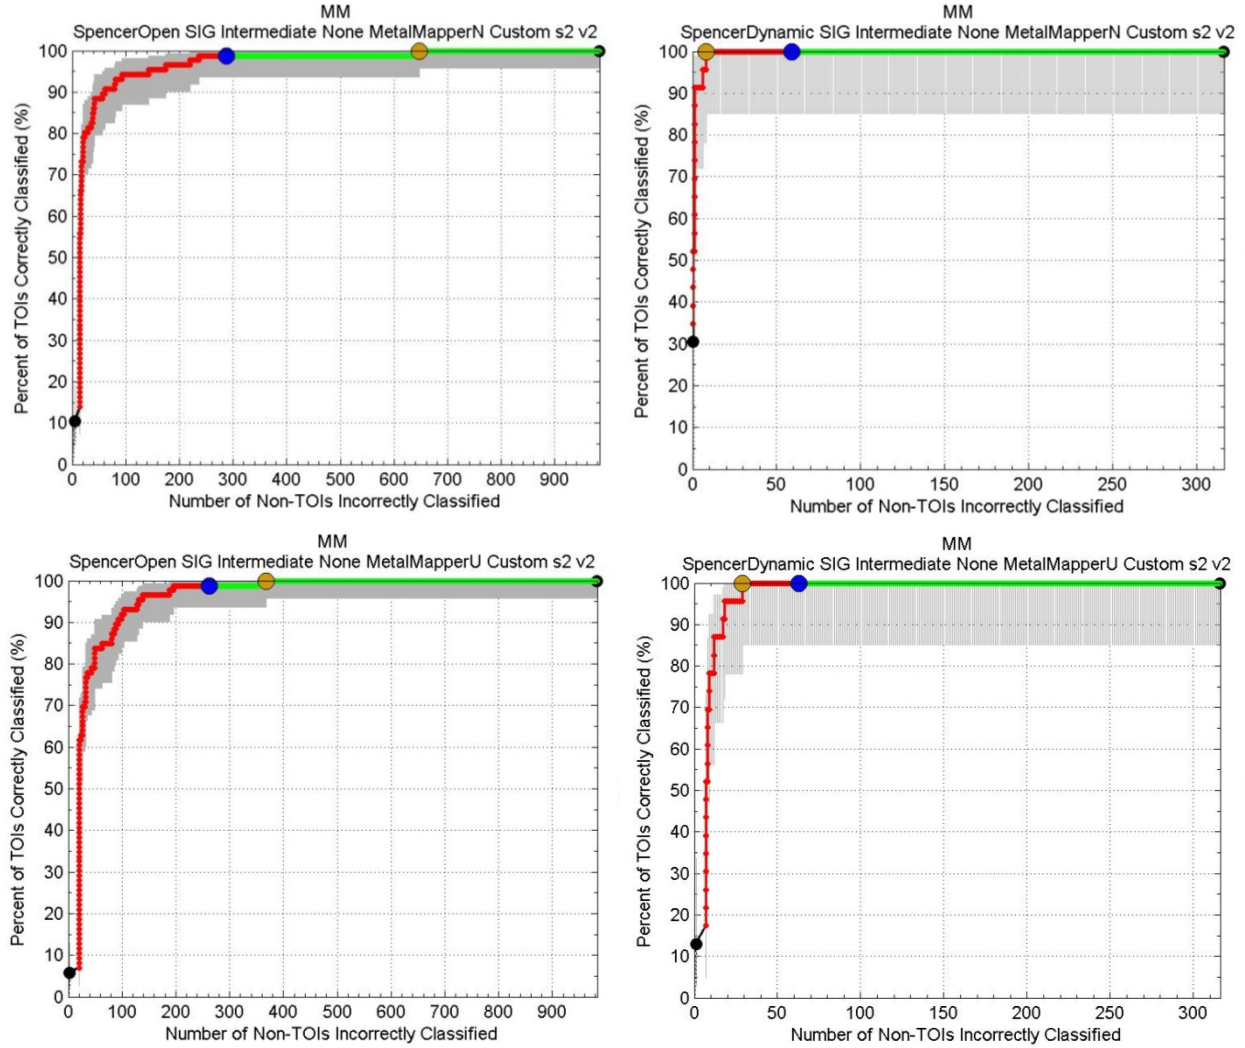


Figure 6. ROC curves for the Intermediate approach, NAEVA and URS sensors(top, bottom), Open and Dynamic areas (left, right).

Having adjusted the stop-dig threshold to accommodate the missed QC seeds that should have modeled as two anomalies and having increased the width of the non-linear kernels, the final dig list for the Dynamic area captured all UXOs with 82 digs. The final dig list for the Open area increased to 387 digs missing one UXO (Figure 6, SR-181). Clutter was dug in the Dynamic area well beyond the last UXO. This was due to the fact that the stop-dig threshold had to be adjusted to accommodate the two-anomaly missed seeds.

URS MetalMapper – Open and Dynamic Areas

The initial dig list contained 233 digs for the Open site and 55 digs for the Dynamic site. These initial dig lists missed eight Q.C. seeds (SR-190, SR-199, SR-837, SR-873, SR-886, SR-1502, SR-1609, SR-1676). As with the NAEVA datasets, most of the missed seeds were attributable to either multi-anomaly responses that had good one-anomaly fits in SIGs inversion model, or were small ISO that were not represented in the test pit data. SR-1502- and SR-1609, however, should have been labeled as ‘can’t extract reliable features’. As a result, the error threshold of the

inversion model which determines whether or not an anomaly is deemed to have good features had to be adjusted. Thus, the final dig lists for *all* modeling method that used the URS data included a few extra anomalies labeled ‘can’t extract reliable features’. The final dig lists for the intermediate approach on the URS dataset had 363 digs for the Open site and 86 digs for the Dynamic site. One UXO, SR-194, was missed in the Open Area.

Conservative Approach

The conservative approach matches the previous classification methods that were performed by SIG at sites such as Camp Beale and Camp Butner. More training data were collected using the conservative method than either the intermediate or aggressive approaches. Training began with a selection of 10 anomalies based on their cumulative Fisher Information. From this set, a PNBC classifier was trained, and active learning was applied to generate a candidate list of new training anomalies. Feature selection for this and subsequent classifiers was performed using the BENet. The feature set converged for all conservative classifier relatively quickly to include the magnitude and decay rates of the first polarizability axis, a measure of symmetry between the 2nd and 3rd axes, and eccentricity. Once additional training labels were acquired, the process was repeated. Training stopped when the information gain of new training points was small relative to information contained in all the previously acquired points. A dig list based on the training data was then generated.

NAEVA MetalMapper – Open and Dynamic Areas

Three rounds of training were performed for a total of 37 labels acquired between the Open and Dynamic Areas. The initial dig list contained 249 digs for the Open site and 56 digs for the Dynamic site. These initial dig lists missed eight QC seeds (SR-490, SR-572, SR-648, SR-790, SR-873, SR-978, SR-1502, SR-1609). All of the missed seed were attributable to either multi-anomaly responses that had good one-anomaly fits in SIGs inversion model, or a maladjustment of the threshold for determining whether anomalies had reliable features or not. The former issue was addressed by changing the stop-dig threshold, and the latter was addressed by labeling more anomalies as ‘can’t extract reliable features’.

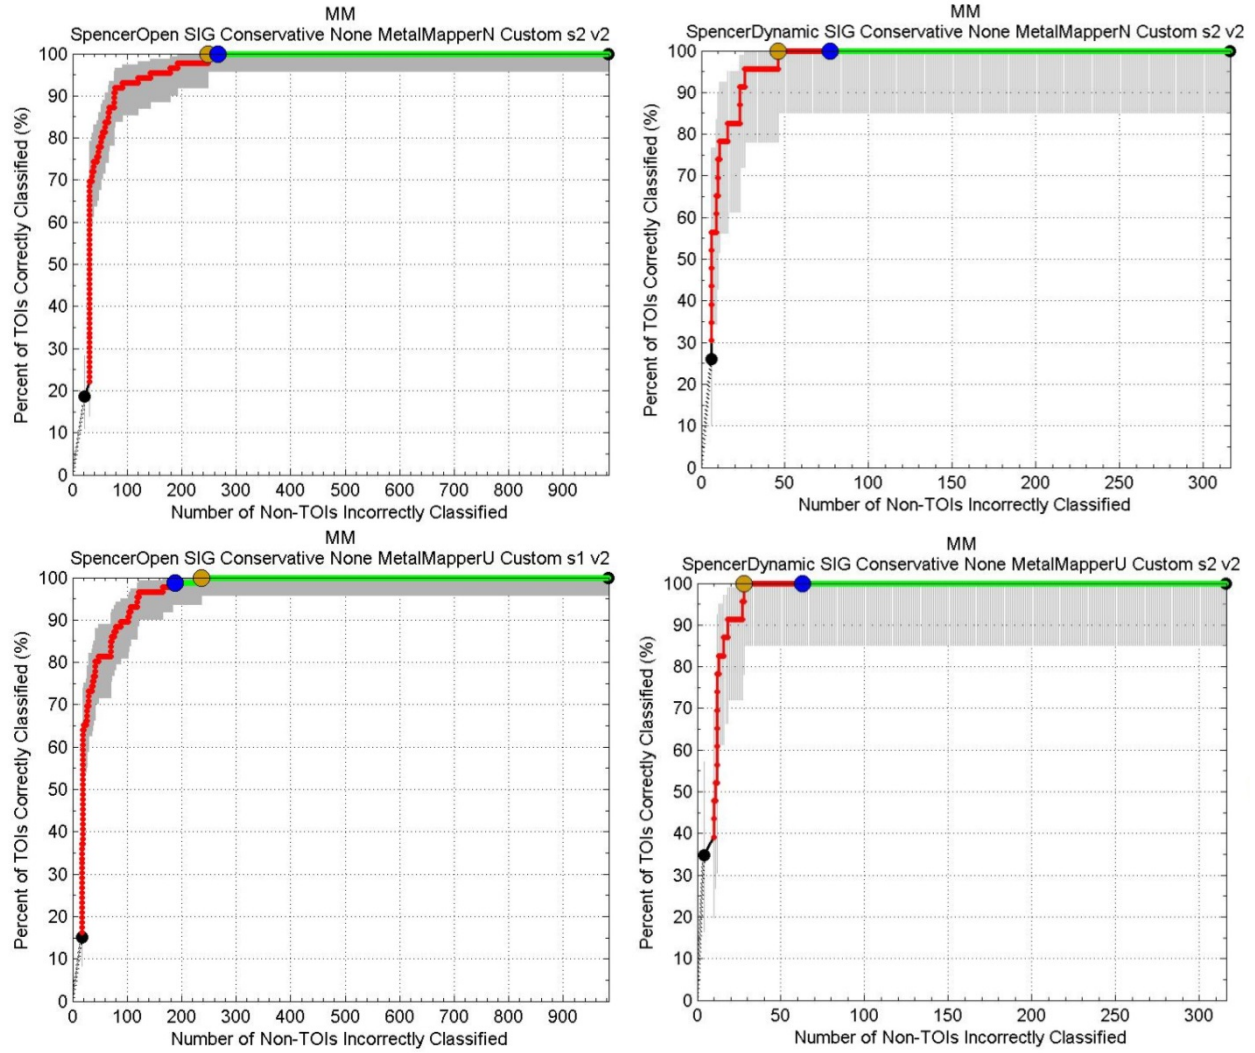


Figure 7. ROC curves for the Conservative approach, NAEVA and URS sensors(top, bottom), Open and Dynamic areas (left, right).

Having adjusted the stop-dig threshold to accommodate the missed QC seeds that were two

anomalies and having increased the number of labels for ‘cant’ extract reliable features’, the final dig list for the Dynamic area captured all UXOs with 100 digs and the final dig list for the Open area increased to 369 digs capturing all UXO (Figure 7). Clutter was dug in the Dynamic area well beyond the last UXO. As with the intermediate approach, this was due to the fact that the stop-dig threshold had to be adjusted to accommodate the two-anomaly missed seeds.

URS MetalMapper – Open and Dynamic Areas

Three rounds of training were performed for a total of 37 labels acquired between the Open and Dynamic areas. The initial dig list contained 249 digs for the Open site and 56 digs for the Dynamic site. These initial dig lists missed two QC seeds (SR-1502, SR-1609). As with the NAEVA dataset all of the missed seed were attributable to either multi-anomaly responses that had good one-anomaly fits in SIGs inversion model, or a maladjustment of the threshold for determining whether anomalies had reliable features or not. The former issue was addressed by changing the stop-dig threshold, and the latter was addressed by labeling more anomalies as ‘can’t extract reliable features’.

Having adjusted the stop-dig threshold to accommodate the missed QC seeds that were two anomalies and having increased the number of labels for ‘cant’ extract reliable features’, the final dig list for the Dynamic area captured all UXOs with 86 digs and the final dig list for the Open area increased to 284 digs missing one-anomaly (Figure 7, SR-194). Both of the missed anomalies were of the Small ISO type not represented in the test pit data. Clutter was dug in the Dynamic area well beyond the last UXO. As with the intermediate approach, this was due to the fact that the stop-dig threshold had to be adjusted to accommodate the two-anomaly missed seeds.

UXO Type Assignment

	0	37	38	48	60	61	75	105	Total
0	0	282	25	0	3	0	6	0	316
37	0	3	4	0	0	0	0	0	7
38	0	0	0	0	0	0	0	0	0
48	0	0	6	0	0	0	0	0	6
60	0	0	1	0	0	0	0	0	1
61	0	0	0	0	1	0	0	0	1
75	0	0	0	0	0	0	4	0	4
105	0	0	0	0	0	0	0	1	1
Total	0	285	36	0	4	0	10	1	336

0 = Non-TOI
 37 = 37mm
 38 = 38mm
 48 = Small ISO
 60 = 60mm
 61 = Medium ISO
 75 = 75mm
 105 = 105mm

Figure 9. UXO type cross-tabulation for the Dynamic site.

Assigning UXO type to anomalies predicted to be UXO was a new requirement for the dig lists. Some of SIGs modeling approaches (intermediate and conservative) treated all UXO as a single type. As a result a single model was used to assign UXO type from a library of UXO responses from previous sites. The features for this model were the first time gate response and the fitted magnitude of axes 2 and 3. So, UXO Type was identical for each dig lists. For the Dynamic Area a single UXO was given an inappropriate classification (Figure 9, i.e. one where the actual diameter was much larger than the predicted diameter). 3 were given inappropriate classifications in the Open site (Figure 8).

Missed UXO

The final number of SIG dig lists was six for the Dynamic area and six for the Open area. No anomalies were missed in the Dynamic area dig list, and in fact many of the dig list dug clutter were beyond the last UXO. Five UXO were missed in the Open dig lists. These five can be separated by sensor. Three were unique to the NAEVA dataset, and two were unique to the URS dataset. All these missed UXO can be attributed to poor features for the dataset in question. So, for the NAEVA missed UXO, the URS features of the same anomalies were easily discriminated as UXO. The same is true for the URS data.

	0	37	38	48	57	60	61	75	105	155
0	716	207	14	0	1	2	0	10	2	0
37	0	11	11	0	0	1	0	1	0	0
38	0	0	0	0	0	0	0	0	0	0
48	0	1	21	0	0	0	0	0	0	0
57	0	0	0	0	0	0	0	0	0	0
60	0	0	3	0	0	1	0	0	0	0
61	0	0	0	0	0	2	0	1	0	0
75	0	0	0	0	0	0	0	14	0	0
105	0	0	0	0	0	0	0	0	0	0
155	0	0	0	0	0	0	0	0	1	0

Figure 8. UXO type cross-tabulation for the Open site.




Figure 10. Polarizabilities for missed UXO in the NAEVA dig lists. Responses for NAEVA (left) and URS (right) are both shown.

Flagged anomalies SR-181, SR-383, and SR-710 were missed in the NAEVA dataset, aggressive

approach (Figure 10). SR-181 was missed in the NAEVA dataset, intermediate approach, and no UXO were missed in the NAEVA dataset, conservative approach. The inversion model for SR-181 had the 2nd and 3rd axes confused. So, the decay rate of the 2nd axis was much lower than expected for a UXO. Since the decay rate of the 2nd axis was selected as an important feature by the NAEVA intermediate and aggressive classifiers, this switch of axes had a significant impact on the probability of UXO for SR-181. The magnitude of all axes for SR-383 was greater than expected for that UXO, a small ISO. Additionally, it was missed in the aggressive approach for which there were no training data of that UXO type. SR-710 was another case like SR-181 where the 2nd and 3rd polarizability axes were reversed relative to what they should have been. It should be noted these anomalies are clearly symmetric in the 2nd and 3rd axes in the NAEVA data, yet they were missed. Symmetry in SIG's methodology is calculated as the sum of squared differences between the log polarizabilities of the 2nd and 3rd axes. The learned weight of the symmetry feature for the NAEVA aggressive approach and intermediate approaches was small relative to the weights placed on the magnitude and decay rate of the 2nd axis. In the case of the missed anomalies, the magnitude of the 2nd axis did not match anything in our training sets for the intermediate and aggressive approaches. Therefore they were not classified as UXO. Additionally, symmetry is calculated on log-polarizabilities. So, it is sensitive to noise in the late time gates. This is good, in general, for discriminating UXO from clutter since UXO typically have less noise in the later time gates than clutter. In the case of the missed anomalies from Figure 10, however, this sensitivity was a detriment since the NAEVA data was noisier at later times than the URS data.

UXOs SR-194 and SR-633 were missed in the aggressive, URS, Open area dig list (Figure 11). SR-194 was missed in the intermediate and conservative approaches. This UXO was another example, along with many of the missed QC seeds where the one-anomaly model fit was extremely low, suggesting that this anomaly was a one-anomaly clutter. But, in fact, it is a two-anomaly UXO. The stop-dig threshold was changed to accommodate these cases for the QC seeds, but not enough to capture SR-194. SR-633 was a small ISO. It was missed in the aggressive approach primarily because there were no examples of these in the test pit data on which the aggressive approach was based.

Comparison between aggressive, intermediate and conservative approaches.

Table 2 summarizes the performance results between the three different approaches: aggressive, intermediate and conservative. The aggressive approach had the fewest digs overall, but also missed the most UXO. The UXO that were missed, however, tended either be poor inversions or more importantly were of UXO types that were not present in the test pit data. This was true for both the missed QC seeds and the UXO missed in the final dig lists. In particular, the UXO labeled small ISO (which was schedule 80 pipe) SIG had no representative in any of the previous sites for the small ISO present in the test data. This result confirms expectations that the




Figure 11. Polarizabilities for missed UXO in the URS dig lists. Responses for NAEVA (left) and URS (right) are both shown.

aggressive approach should work well when two conditions are met. The first condition is that sites have relatively low feature complexity. Camp Spencer was intermediately complex (See ‘Evaluation of site complexity’) so expectations on the aggressive approach were low. The second condition for the aggressive approach is that a generative model exists for each UXO type in the test data. There was no representative for the small ISO. Further, the URS dataset had only 3 UXO types in its test pit data, while the NAEVA data had 6. A consequence of the increased diversity of test pit data was that classification on the NAEVA dataset dug UXO with fewer false alarms than classifications of the URS data. This highlights both the strengths and weakness of generative models and the similar library matching approach. UXO can be dug with very few false alarms, but these models are poor at exploring the feature space for UXO that are not part of the library.

The intermediate approach required more digs than any other approach. But, it was also more strongly affected by the misclassification of two-anomaly responses. There were few training labels, and by design all of them were near the boundary between UXO and clutter in feature space. In the context of a non-linear classifier, this meant that each clutter dug had a large relative importance. The two-anomaly missed QC seeds were far away from the classifier boundary defined by the training points. So, when the stop-dig threshold was adjusted to accommodate the two-anomaly missed QC seeds, many more clutter dug. There just were not enough training labels in the area near the missed QC seeds to differentiate any additional misclassified two-anomaly responses or to reduce the uncertainty between UXO and clutter in that area of feature space.

The conservative approach was the only approach to capture all UXO in both sites using the NAEVA dataset. However, this list also had the largest number of false positives and the largest number of missed QC seeds. The conservative approach applied to the URS dataset had many fewer false positives or missed QC seeds than the NAEVA dataset. This classification did, however, miss a single UXO. The performance was equivalent to the intermediate approach with approximately 30% fewer false positives. The NAEVA discrimination benefitted from having twice as many of the misclassified two-anomaly responses present in the training dataset. This illustrates the benefit of using the active learning of the conservative approach in cases where little is known about a site in terms of the types of UXO present. Though some the two-anomaly responses were ‘misclassified’ as one-anomaly responses, the conservative approach

	Open				Dynamic			
	can't analyze	training	dug	missed UXO	can't analyze	training	labeld	missed UXO
Aggr.	3	0 (0)	210	3	0	0 (2)	57	0
Naeva Inter.	11	9 (6)	361	1	0	4 (3)	75	0
Cons.	11	32 (6)	320	0	0	10 (12)	88	0
Aggr.	19	0 (3)	278	2	7	0 (2)	62	0
URS Inter.	19	2 (5)	337	1	6	1 (3)	76	0
Cons.	1	37 (0)	253	1	6	10 (2)	68	0

Table 2. Summary of the performance for the Aggressive (Aggr.), Intermediate (Inter.), and Conservative (Cons.) approaches. Values are counts of flagged anomalies. For training data, missed seeds are shown in parentheses since they were required to be part of the training data but were not part of the training in the classification models.

was still able to capture them using active learning. Further, once a single instance was captured, the active learning delineates an effective classifier boundary around these types of responses even though they were being treated as one-anomaly responses. This confirms the expectation that the conservative approach works the best when the site is complex and little information is available about the types of UXO present in the test data.

Evaluation of site complexity

One goal of the Camp Spencer discrimination was to develop a method for determining *a priori* which of the three modeling approaches to use at a given site. SIG computed a metric of site complexity for the Camp Spencer MetalMapper datasets as well as former sites Pole Mountain, Camp Butner, and Camp Beale MetalMapper datasets. Retrospective analysis of these sites had suggested that the Pole Mountain Maneuver Area was easy to discriminate. The aggressive approach applied to that site outperformed the conservative approach, capturing all the UXO with far fewer false positives. Camp Beale and Camp Butner were more difficult. In both these cases the conservative approach performed best. Performance of the different modeling approaches at Camp Spencer would suggest that conservative or intermediate approaches should have performed best. The intermediate approach, in particular should have performed much better had there been some representatives of the small ISO in the test pit data and if the two-anomaly responses had been classified better.

Table 3. Complexity metric for the current site and three former sites. Larger values suggest more intrinsic complexity and a more conservative modeling approach.

Site	Complexity
Camp Spencer, URS	0.1927
Camp Spencer, Naeva	0.1931
Pole Mountain, MetalMapper	0.1705
Camp Beale, MetalMapper	0.2019
Camp Butner, MetalMapper	0.2075

The complexity metric calculated for all these sites is shown in Table 3. It should be noted that while the support of the complexity metric is (0, 1) it is unlikely that the complexity of any site will approach 1. A value near 1 would suggest there is little to no random variation associated with sensor, site, clutter properties or UXO properties and that the separation between UXO types and clutter spans the entire width of the feature space. This scenario is unlikely, and such a site would not need to be discriminated using a statistical approach. Thus, one must assess the relative complexity among sites not the value of the complexity metric as compared to its theoretical support.

This complexity metric calculated at the different sites exhibits a strong correlation with the observed difficulty of the sites. While there are too few sites to make a definite cutoff for the type of model to use given the complexity metric some general rules might be suggested. Values near 0.2 should be conservative. Values at or near below 0.17 should be aggressive. But, as the Camp Spencer demonstration illustrates, the measure of site complexity must not be the only factor in choosing a particular model. Confidence about the expected UXO types must also be considered.

Conclusions and Implications for Future Research/Implementation

The SIG discrimination process has proven effective at providing efficient site discrimination with minimal training data. Former Spencer Artillery Range provides another example of this. The addition of adaptive modeling approaches that use an *a priori* metric of site complexity has increased effectiveness of this approach. The aggressive approach has the capacity to capture UXO with the least expense – if the site is of low complexity and there is a large diversity of UXO represented in the test pit data. The conservative approach is best when little is known about a site or there are sources of noise such as high-iron soils that make the discrimination more complicated. The intermediate approach can be used for sites where UXO types are well known, but the site itself is too complex for the aggressive approach. An area of improvement for the SIG discrimination process is defining a better method for deciding whether a given flag should have the one-anomaly applied or a multi-anomaly model applied. Some of the missed QC seeds that were two-anomaly responses had one-anomaly fit errors that were lower than 95% of the other one-anomaly responses. A more refined error metric or a method for simultaneously including both the one-anomaly responses and the two-anomaly responses should be examined.

Literature Cited

- [1] Y. Zhang, X. Liao, and L. Carin, “Detection of buried targets via active selection of labeled data: application to sensing subsurface UXO,” IEEE Transactions on Geoscience and Remote Sensing, vol. 42, pp. 2535–2543, November 2004.
- [2] Y. Zhang, H. Y. L. Collins, C. Baum, and L. Carin, “Sensing of unexploded ordnance with magnetometer and induction data: Theory and signal processing,” IEEE Transactions on Geoscience and Remote Sensing, vol. 41, pp. 1005–1015, May 2003.
- [3] N. Geng, C. E. Baum, and L. Carin, “On the low-frequency natural response of conducting and permeable target,” IEEE Transactions on Geoscience and Remote Sensing, vol. 37, pp. 347–359, January 1999.
- [4] L. Carin, H. Yu, Y. Dalichaouch, A. Perry, and P. Czipott, “On the wideband EMI response of a rotationally symmetric permeable and conducting target,” IEEE Transactions on Geoscience and Remote Sensing, 2001.
- [5] Y. Zhang, X. Liao, and L. Carin, “Detection of buried targets via active selection of labeled data: application to sensing subsurface UXO,” IEEE Transactions on Geoscience and Remote Sensing, vol. 42, pp. 2535–2543, November 2004.
- [6] Kowalski, Andres M., et al. "Distances in probability space and the statistical complexity setup." Entropy. Vol. 13, pp. 1055-1075. 2011
- [7] N. Cristianini and J. Shawe-Taylor, An Introduction to Support Vector Machines and other kernel-based learning methods. Cambridge University Press, 2000.
- [8] X. Liao and L. Carin, “Migratory Logistic Regression for Learning Concept Drift Between Two Data Sets With Application to UXO Sensing,” IEEE Transactions on Geoscience and Remote Sensing, vol. 46, no. 12, December 2008.

[9] D. MacKay, “Information-based objective functions for active data selection,” *Neural Computation*, vol. 4, pp. 589–603, 1992.

Appendices

Appendix A: Points of Contact

POINT OF CONTACT Name	ORGANIZATION Name Address	Phone Fax E-mail	Role in Project
Lawrence Carin	Signal Innovations Group, Inc. 4721 Emperor Blvd., Suite 330 Durham, NC 27703	919 660-5270 919-323-4811 lcarin@ece.duke.edu	Principal Investigator
Levi Kennedy	Signal Innovations Group, Inc. 4721 Emperor Blvd., Suite 330 Durham, NC 27703	919-323-3456 919-287-2578 lkennedy@siginnovations.com	Project Management
Todd Jobe	Signal Innovations Group, Inc. 4721 Emperor Blvd., Suite 330 Durham, NC 27703	919-323-4811 919-287-2578 tjobe@siginnovations.com	Engineer